Springer Finance

Steven E. Shreve

Stochastic Calculus for Finance I

The Binomial Asset Pricing Model

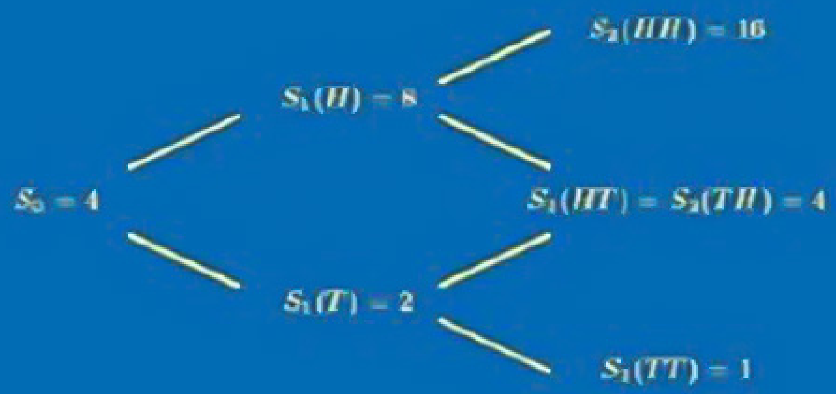

Springer

Steven Shreve: Stochastic Calculus and Finance

PRASAD CHALASANI
Carnegie Mellon University
chal@cs.cmu.edu

SOMESH JHA
Carnegie Mellon University
sjha@cs.cmu.edu

©Copyright; Steven E. Shreve, 1996

July 25, 1997

Contents

Chapter 1

Introduction to Probability Theory

1.1 The Binomial Asset Pricing Model

The *binomial asset pricing model* provides a powerful tool to understand arbitrage pricing theory and probability theory. In this course, we shall use it for both these purposes.

In the binomial asset pricing model, we model stock prices in discrete time, assuming that at each step, the stock price will change to one of two possible values. Let us begin with an initial positive stock price S_0. There are two positive numbers, d and u, with

$$0 < d < u, \tag{1.1}$$

such that at the next period, the stock price will be either dS_0 or uS_0. Typically, we take d and u to satisfy $0 < d < 1 < u$, so change of the stock price from S_0 to dS_0 represents a *downward* movement, and change of the stock price from S_0 to uS_0 represents an *upward* movement. It is common to also have $d = \frac{1}{u}$, and this will be the case in many of our examples. However, strictly speaking, for what we are about to do we need to assume only (1.1) and (1.2) below.

Of course, stock price movements are much more complicated than indicated by the binomial asset pricing model. We consider this simple model for three reasons. First of all, within this model the concept of arbitrage pricing and its relation to risk-neutral pricing is clearly illuminated. Secondly, the model is used in practice because with a sufficient number of steps, it provides a good, computationally tractable approximation to continuous-time models. Thirdly, within the binomial model we can develop the theory of conditional expectations and martingales which lies at the heart of continuous-time models.

With this third motivation in mind, we develop notation for the binomial model which is a bit different from that normally found in practice. Let us imagine that we are tossing a coin, and when we get a "Head," the stock price moves up, but when we get a "Tail," the price moves down. We denote the price at time 1 by $S_1(H) = uS_0$ if the toss results in head (H), and by $S_1(T) = dS_0$ if it

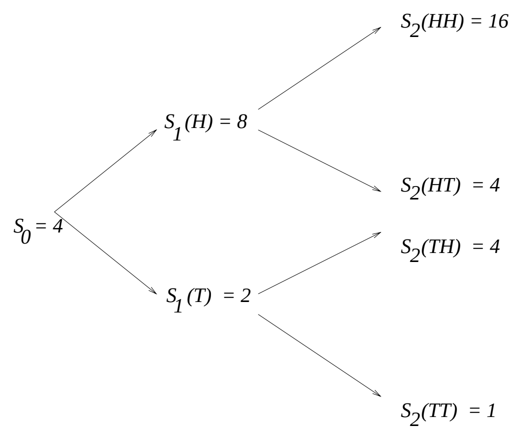

Figure 1.1: *Binomial tree of stock prices with* $S_0 = 4$*,* $u = 1/d = 2$*.*

results in tail (T). After the second toss, the price will be one of:

$$S_2(HH) = uS_1(H) = u^2 S_0, \quad S_2(HT) = dS_1(H) = duS_0,$$

$$S_2(TH) = uS_1(T) = udS_0, \quad S_2(TT) = dS_1(T) = d^2 S_0.$$

After three tosses, there are eight possible coin sequences, although not all of them result in different stock prices at time 3.

For the moment, let us assume that the third toss is the last one and denote by

$$\Omega = \{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT\}$$

the set of all possible outcomes of the three tosses. The set Ω of all possible outcomes of a random experiment is called the *sample space* for the experiment, and the elements ω of Ω are called *sample points*. In this case, each sample point ω is a sequence of length three. We denote the k-th component of ω by ω_k. For example, when $\omega = HTH$, we have $\omega_1 = H$, $\omega_2 = T$ and $\omega_3 = H$.

The stock price S_k at time k depends on the coin tosses. To emphasize this, we often write $S_k(\omega)$. Actually, this notation does not quite tell the whole story, for while S_3 depends on all of ω, S_2 depends on only the first two components of ω, S_1 depends on only the first component of ω, and S_0 does not depend on ω at all. Sometimes we will use notation such $S_2(\omega_1, \omega_2)$ just to record more explicitly how S_2 depends on $\omega = (\omega_1, \omega_2, \omega_3)$.

Example 1.1 Set $S_0 = 4$, $u = 2$ and $d = \frac{1}{2}$. We have then the binomial "tree" of possible stock prices shown in Fig. 1.1. Each sample point $\omega = (\omega_1, \omega_2, \omega_3)$ represents a path through the tree. Thus, we can think of the sample space Ω as either the set of all possible outcomes from three coin tosses or as the set of all possible paths through the tree.

To complete our binomial asset pricing model, we introduce a *money market* with *interest rate* r; \$1 invested in the money market becomes $\$(1 + r)$ in the next period. We take r to be the interest rate for both *borrowing* and *lending*. (This is not as ridiculous as it first seems, because in a many applications of the model, an agent is either borrowing or lending (not both) and knows in advance which she will be doing; in such an application, she should take r to be the rate of interest for her activity.) We assume that

$$d < 1 + r < u. \tag{1.2}$$

The model would not make sense if we did not have this condition. For example, if $1 + r \geq u$, then the rate of return on the money market is always at least as great as and sometimes greater than the return on the stock, and no one would invest in the stock. The inequality $d \geq 1 + r$ cannot happen unless either r is negative (which never happens, except maybe once upon a time in Switzerland) or $d \geq 1$. In the latter case, the stock does not really go "down" if we get a tail; it just goes up less than if we had gotten a head. One should borrow money at interest rate r and invest in the stock, since even in the worst case, the stock price rises at least as fast as the debt used to buy it.

With the stock as the underlying asset, let us consider a *European call option* with strike price $K > 0$ and expiration time 1. This option confers the right to buy the stock at time 1 for K dollars, and so is worth $S_1 - K$ at time 1 if $S_1 - K$ is positive and is otherwise worth zero. We denote by

$$V_1(\omega) = (S_1(\omega) - K)^+ \stackrel{\Delta}{=} \max\{S_1(\omega) - K, 0\}$$

the value (payoff) of this option at expiration. Of course, $V_1(\omega)$ actually depends only on ω_1, and we can and do sometimes write $V_1(\omega_1)$ rather than $V_1(\omega)$. Our first task is to compute the *arbitrage price* of this option at time zero.

Suppose at time zero you sell the call for V_0 dollars, where V_0 is still to be determined. You now have an obligation to pay off $(uS_0 - K)^+$ if $\omega_1 = H$ and to pay off $(dS_0 - K)^+$ if $\omega_1 = T$. At the time you sell the option, you don't yet know which value ω_1 will take. You *hedge* your short position in the option by buying Δ_0 shares of stock, where Δ_0 is still to be determined. You can use the proceeds V_0 of the sale of the option for this purpose, and then borrow if necessary at interest rate r to complete the purchase. If V_0 is more than necessary to buy the Δ_0 shares of stock, you invest the residual money at interest rate r. In either case, you will have $V_0 - \Delta_0 S_0$ dollars invested in the money market, where this quantity might be negative. You will also own Δ_0 shares of stock.

If the stock goes up, the value of your portfolio (excluding the short position in the option) is

$$\Delta_0 S_1(H) + (1 + r)(V_0 - \Delta_0 S_0),$$

and you need to have $V_1(H)$. Thus, you want to choose V_0 and Δ_0 so that

$$V_1(H) = \Delta_0 S_1(H) + (1 + r)(V_0 - \Delta_0 S_0). \tag{1.3}$$

If the stock goes down, the value of your portfolio is

$$\Delta_0 S_1(T) + (1 + r)(V_0 - \Delta_0 S_0),$$

and you need to have $V_1(T)$. Thus, you want to choose V_0 and Δ_0 to also have

$$V_1(T) = \Delta_0 S_1(T) + (1 + r)(V_0 - \Delta_0 S_0). \tag{1.4}$$

These are two equations in two unknowns, and we solve them below

Subtracting (1.4) from (1.3), we obtain

$$V_1(H) - V_1(T) = \Delta_0(S_1(H) - S_1(T)), \tag{1.5}$$

so that

$$\Delta_0 = \frac{V_1(H) - V_1(T)}{S_1(H) - S_1(T)}. \tag{1.6}$$

This is a discrete-time version of the famous "delta-hedging" formula for derivative securities, according to which the number of shares of an underlying asset a hedge should hold is the derivative (in the sense of calculus) of the value of the derivative security with respect to the price of the underlying asset. This formula is so pervasive the when a practitioner says "delta", she means the derivative (in the sense of calculus) just described. Note, however, that my *definition* of Δ_0 is the number of shares of stock one holds at time zero, and (1.6) is a consequence of this definition, not the definition of Δ_0 itself. Depending on how uncertainty enters the model, there can be cases in which the number of shares of stock a hedge should hold is not the (calculus) derivative of the derivative security with respect to the price of the underlying asset.

To complete the solution of (1.3) and (1.4), we substitute (1.6) into either (1.3) or (1.4) and solve for V_0. After some simplification, this leads to the formula

$$V_0 = \frac{1}{1+r}\left[\frac{1+r-d}{u-d}V_1(H) + \frac{u-(1+r)}{u-d}V_1(T)\right]. \tag{1.7}$$

This is the *arbitrage price* for the European call option with payoff V_1 at time 1. To simplify this formula, we define

$$\tilde{p} \triangleq \frac{1+r-d}{u-d}, \quad \tilde{q} \triangleq \frac{u-(1+r)}{u-d} = 1 - \tilde{p}, \tag{1.8}$$

so that (1.7) becomes

$$V_0 = \frac{1}{1+r}[\tilde{p}V_1(H) + \tilde{q}V_1(T)]. \tag{1.9}$$

Because we have taken $d < u$, both $\tilde{p}$ and $\tilde{q}$ are defined,i.e., the denominator in (1.8) is not zero. Because of (1.2), both $\tilde{p}$ and $\tilde{q}$ are in the interval $(0,1)$, and because they sum to 1, we can regard them as probabilities of H and T, respectively. They are the *risk-neutral* probabilites. They appeared when we solved the two equations (1.3) and (1.4), and have nothing to do with the actual probabilities of getting H or T on the coin tosses. In fact, at this point, they are nothing more than a convenient tool for writing (1.7) as (1.9).

We now consider a European call which pays off K dollars at time 2. At expiration, the payoff of this option is $V_2 \triangleq (S_2 - K)^+$, where V_2 and S_2 depend on ω_1 and ω_2, the first and second coin tosses. We want to determine the arbitrage price for this option at time zero. Suppose an agent sells the option at time zero for V_0 dollars, where V_0 is still to be determined. She then buys Δ_0 shares of stock, investing $V_0 - \Delta_0 S_0$ dollars in the money market to finance this. At time 1, the agent has a portfolio (excluding the short position in the option) valued at

$$X_1 \triangleq \Delta_0 S_1 + (1+r)(V_0 - \Delta_0 S_0). \tag{1.10}$$

Although we do not indicate it in the notation, S_1 and therefore X_1 depend on ω_1, the outcome of the first coin toss. Thus, there are really two equations implicit in (1.10):

$$\begin{aligned} X_1(H) &\triangleq \Delta_0 S_1(H) + (1+r)(V_0 - \Delta_0 S_0), \\ X_1(T) &\triangleq \Delta_0 S_1(T) + (1+r)(V_0 - \Delta_0 S_0). \end{aligned}$$

After the first coin toss, the agent has X_1 dollars and can readjust her hedge. Suppose she decides to now hold Δ_1 shares of stock, where Δ_1 is allowed to depend on ω_1 because the agent knows what value ω_1 has taken. She invests the remainder of her wealth, $X_1 - \Delta_1 S_1$ in the money market. In the next period, her wealth will be given by the right-hand side of the following equation, and she wants it to be V_2. Therefore, she wants to have

$$V_2 = \Delta_1 S_2 + (1+r)(X_1 - \Delta_1 S_1). \tag{1.11}$$

Although we do not indicate it in the notation, S_2 and V_2 depend on ω_1 and ω_2, the outcomes of the first two coin tosses. Considering all four possible outcomes, we can write (1.11) as four equations:

$$\begin{aligned} V_2(HH) &= \Delta_1(H)S_2(HH) + (1+r)(X_1(H) - \Delta_1(H)S_1(H)), \\ V_2(HT) &= \Delta_1(H)S_2(HT) + (1+r)(X_1(H) - \Delta_1(H)S_1(H)), \\ V_2(TH) &= \Delta_1(T)S_2(TH) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)), \\ V_2(TT) &= \Delta_1(T)S_2(TT) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)). \end{aligned}$$

We now have six equations, the two represented by (1.10) and the four represented by (1.11), in the six unknowns V_0, Δ_0, $\Delta_1(H)$, $\Delta_1(T)$, $X_1(H)$, and $X_1(T)$.

To solve these equations, and thereby determine the arbitrage price V_0 at time zero of the option and the hedging portfolio Δ_0, $\Delta_1(H)$ and $\Delta_1(T)$, we begin with the last two

$$\begin{aligned} V_2(TH) &= \Delta_1(T)S_2(TH) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)), \\ V_2(TT) &= \Delta_1(T)S_2(TT) + (1+r)(X_1(T) - \Delta_1(T)S_1(T)). \end{aligned}$$

Subtracting one of these from the other and solving for $\Delta_1(T)$, we obtain the "delta-hedging formula"

$$\Delta_1(T) = \frac{V_2(TH) - V_2(TT)}{S_2(TH) - S_2(TT)}, \tag{1.12}$$

and substituting this into either equation, we can solve for

$$X_1(T) = \frac{1}{1+r}[\tilde{p}V_2(TH) + \tilde{q}V_2(TT)]. \tag{1.13}$$

Equation (1.13), gives the value the hedging portfolio should have at time 1 if the stock goes down between times 0 and 1. We define this quantity to be the *arbitrage value of the option at time 1 if $\omega_1 = T$*, and we denote it by $V_1(T)$. We have just shown that

$$V_1(T) \triangleq \frac{1}{1+r}[\tilde{p}V_2(TH) + \tilde{q}V_2(TT)]. \tag{1.14}$$

The hedger should choose her portfolio so that her wealth $X_1(T)$ if $\omega_1 = T$ agrees with $V_1(T)$ defined by (1.14). This formula is analgous to formula (1.9), but postponed by one step. The first two equations implicit in (1.11) lead in a similar way to the formulas

$$\Delta_1(H) = \frac{V_2(HH) - V_2(HT)}{S_2(HH) - S_2(HT)} \tag{1.15}$$

and $X_1(H) = V_1(H)$, where $V_1(H)$ is the value of the option at time 1 if $\omega_1 = H$, defined by

$$V_1(H) \triangleq \frac{1}{1+r}[\tilde{p}V_2(HH) + \tilde{q}V_2(HT)]. \tag{1.16}$$

This is again analgous to formula (1.9), postponed by one step. Finally, we plug the values $X_1(H) = V_1(H)$ and $X_1(T) = V_1(T)$ into the two equations implicit in (1.10). The solution of these equations for Δ_0 and V_0 is the same as the solution of (1.3) and (1.4), and results again in (1.6) and (1.9).

The pattern emerging here persists, regardless of the number of periods. If V_k denotes the value at time k of a derivative security, and this depends on the first k coin tosses $\omega_1, \ldots, \omega_k$, then at time $k-1$, after the first $k-1$ tosses $\omega_1, \ldots, \omega_{k-1}$ are known, the portfolio to hedge a short position should hold $\Delta_{k-1}(\omega_1, \ldots, \omega_{k-1})$ shares of stock, where

$$\Delta_{k-1}(\omega_1, \ldots, \omega_{k-1}) = \frac{V_k(\omega_1, \ldots, \omega_{k-1}, H) - V_k(\omega_1, \ldots, \omega_{k-1}, T)}{S_k(\omega_1, \ldots, \omega_{k-1}, H) - S_k(\omega_1, \ldots, \omega_{k-1}, T)}, \tag{1.17}$$

and the value at time $k-1$ of the derivative security, when the first $k-1$ coin tosses result in the outcomes $\omega_1, \ldots, \omega_{k-1}$, is given by

$$V_{k-1}(\omega_1, \ldots, \omega_{k-1}) = \frac{1}{1+r}[\tilde{p}V_k(\omega_1, \ldots, \omega_{k-1}, H) + \tilde{q}V_k(\omega_1, \ldots, \omega_{k-1}, T)] \tag{1.18}$$

1.2 Finite Probability Spaces

Let Ω be a set with finitely many elements. An example to keep in mind is

$$\Omega = \{HHH, HHT, HTH, HTT, THH, THT, TTH, TTT\} \tag{2.1}$$

of all possible outcomes of three coin tosses. Let $\mathcal{F}$ be the set of all subsets of Ω. Some sets in $\mathcal{F}$ are $\emptyset$, $\{HHH, HHT, HTH, HTT\}$, $\{TTT\}$, and Ω itself. How many sets are there in $\mathcal{F}$?

Definition 1.1 A *probability measure* $I\!P$ is a function mapping $\mathcal{F}$ into $[0,1]$ with the following properties:

(i) $I\!P(\Omega) = 1$,

(ii) If $A_1, A_2, \ldots$ is a sequence of disjoint sets in $\mathcal{F}$, then

$$I\!P\left(\bigcup_{k=1}^{\infty} A_k\right) = \sum_{k=1}^{\infty} I\!P(A_k).$$

Probability measures have the following interpretation. Let A be a subset of $\mathcal{F}$. Imagine that Ω is the set of all possible outcomes of some random experiment. There is a certain probability, between 0 and 1, that when that experiment is performed, the outcome will lie in the set A. We think of $I\!P(A)$ as this probability.

Example 1.2 Suppose a coin has probability $\frac{1}{3}$ for H and $\frac{2}{3}$ for T. For the individual elements of Ω in (2.1), define

$$\begin{aligned}
I\!P\{HHH\} &= \left(\tfrac{1}{3}\right)^3, & I\!P\{HHT\} &= \left(\tfrac{1}{3}\right)^2\left(\tfrac{2}{3}\right),\\
I\!P\{HTH\} &= \left(\tfrac{1}{3}\right)^2\left(\tfrac{2}{3}\right), & I\!P\{HTT\} &= \left(\tfrac{1}{3}\right)\left(\tfrac{2}{3}\right)^2,\\
I\!P\{THH\} &= \left(\tfrac{1}{3}\right)^2\left(\tfrac{1}{3}\right), & I\!P\{THT\} &= \left(\tfrac{1}{3}\right)\left(\tfrac{2}{3}\right)^2,\\
I\!P\{TTH\} &= \left(\tfrac{1}{3}\right)\left(\tfrac{2}{3}\right)^2, & I\!P\{TTT\} &= \left(\tfrac{2}{3}\right)^3.
\end{aligned}$$

For $A \in \mathcal{F}$, we define

$$I\!P(A) = \sum_{\omega \in A} I\!P\{\omega\}. \tag{2.2}$$

For example,

$$I\!P\{HHH, HHT, HTH, HTT\} = \left(\frac{1}{3}\right)^3 + 2\left(\frac{1}{3}\right)^2\left(\frac{2}{3}\right) + \left(\frac{1}{3}\right)\left(\frac{2}{3}\right)^2 = \frac{1}{3},$$

which is another way of saying that the probability of H on the first toss is $\frac{1}{3}$.

As in the above example, it is generally the case that we specify a probability measure on only some of the subsets of Ω and then use property (ii) of Definition 1.1 to determine $I\!P(A)$ for the remaining sets $A \in \mathcal{F}$. In the above example, we specified the probability measure only for the sets containing a single element, and then used Definition 1.1(ii) in the form (2.2) (see Problem 1.4(ii)) to determine $I\!P$ for all the other sets in $\mathcal{F}$.

Definition 1.2 Let Ω be a nonempty set. A σ-algebra is a collection $\mathcal{G}$ of subsets of Ω with the following three properties:

(i) $\emptyset \in \mathcal{G}$,

(ii) If $A \in \mathcal{G}$, then its complement $A^c \in \mathcal{G}$,

(iii) If $A_1, A_2, A_3, \ldots$ is a sequence of sets in $\mathcal{G}$, then $\cup_{k=1}^{\infty} A_k$ is also in $\mathcal{G}$.

Here are some important σ-algebras of subsets of the set Ω in Example 1.2:

$$\begin{aligned}
\mathcal{F}_0 &= \Big\{\emptyset, \Omega\Big\}, \\
\mathcal{F}_1 &= \Big\{\emptyset, \Omega, \{HHH, HHT, HTH, HTT\}, \{THH, THT, TTH, TTT\}\Big\}, \\
\mathcal{F}_2 &= \Big\{\emptyset, \Omega, \{HHH, HHT\}, \{HTH, HTT\}, \{THH, THT\}, \{TTH, TTT\}, \\
&\qquad\qquad \text{and all sets which can be built by taking unions of these}\Big\}, \\
\mathcal{F}_3 &= \mathcal{F} = \text{The set of all subsets of } \Omega.
\end{aligned}$$

To simplify notation a bit, let us define

$$\begin{aligned}
A_H &\triangleq \{HHH, HHT, HTH, HTT\} = \{H \text{ on the first toss}\}, \\
A_T &\triangleq \{THH, THT, TTH, TTT\} = \{T \text{ on the first toss}\},
\end{aligned}$$

so that

$$\mathcal{F}_1 = \{\emptyset, \Omega, A_H, A_T\},$$

and let us define

$$\begin{aligned}
A_{HH} &\triangleq \{HHH, HHT\} = \{HH \text{ on the first two tosses}\}, \\
A_{HT} &\triangleq \{HTH, HTT\} = \{HT \text{ on the first two tosses}\}, \\
A_{TH} &\triangleq \{THH, THT\} = \{TH \text{ on the first two tosses}\}, \\
A_{TT} &\triangleq \{TTH, TTT\} = \{TT \text{ on the first two tosses}\},
\end{aligned}$$

so that

$$\begin{aligned}
\mathcal{F}_2 = \;& \{\emptyset, \Omega, A_{HH}, A_{HT}, A_{TH}, A_{TT}, \\
& A_H, A_T, A_{HH} \cup A_{TH}, A_{HH} \cup A_{TT}, A_{HT} \cup A_{TH}, A_{HT} \cup A_{TT}, \\
& A_{HH}^c, A_{HT}^c, A_{TH}^c, A_{TT}^c\}.
\end{aligned}$$

We interpret σ-algebras as a record of information. Suppose the coin is tossed three times, and you are not told the outcome, but you are told, for every set in $\mathcal{F}_1$ whether or not the outcome is in that set. For example, you would be told that the outcome is not in $\emptyset$ and is in Ω. Moreover, you might be told that the outcome is not in A_H but is in A_T. In effect, you have been told that the first toss was a T, and nothing more. The σ-algebra $\mathcal{F}_1$ is said to contain the "information of the first toss", which is usually called the "information up to time 1". Similarly, $\mathcal{F}_2$ contains the "information of